VLSI ARCHITECTURE DESIGN COURSE

LECTURE #1:

- **Course administration**
- **Purpose of the course**
- **Course outline**
- **Trends in VLSI**

COURSE ADMINISTRATION

- **Sunday - 18:30-20:30 - Room 352 EE *<change time?>***
- **Instructors : Uri Weiser, Ronny Ronen**
- **Teaching Assistant : Konstantin Moiseev**
- **Consultation: Email to: [mkostya@techunix.technion.ac.il]**
- **2 home exercises**
- **Final exam/work**
- **Foils/copies**
- **Web: http://www.ee.technion.ac.il/courses/048853/**

PURPOSE OF THE COURSE:

*"**VLSI Architecture"** with emphasis on High performance MicroProcessors within power envelop*

- **Present VLSI Architecture design Considerations**
  - **State of the Art computer architecture**
  - **Process technology implications**
  - Design process
  - Market requirements
  - Economical considerations
- **Present conceptual/practical aspects of VLSI design**
- **Provide a stage for interactive discussions**

COURSE OUTLINE

1. Introduction

- **Course Introduction**
- **Trends in VLSI (e.g. technologies, microarchitectures, circuits...)**

2. Product requirement

- **Driving Forces (e.g. # of transistors, power, integration, applications)**
- **Technology limitations**
- **VLSI Simulations (introduction)**

COURSE OUTLINE (cont)

3. VLSI generic microprocessor

- **The VLSI impact on the generic Microprocessor RISC/CISC VLSI tradeoffs Performance Metrics/Benchmarks**
- **VLSI Microprocessor system partitioning - concept, trends and status**
- **Instruction supply, execution, data supply**
- **Inter/Intra-Chip Communication bottleneck**
- **Potential for instruction level parallelism**
- **Advanced VLSI MicroArchitecture Techniques (e.g Branch Prediction in VLSI, VLSI aspects of Out Of Order...)**

4. Core (execution)

- **Scheduling**
- **Registers**
- **Execution units**
- **OOO example and enhancements**

COURSE OUTLINE (cont)

5. **Instruction supply**
   - **Decoding**
   - **Icache, trace cache**
6. **Instruction supply (cont)**
   - **Branch prediction**
   - **Speculative execution**
7. **Data Supply**
   - **Caches**
   - **Memory data prefetch**
   - **Data speculation**
8. **Data Supply (cont)**
   - **Memory disambiguation**
   - **On die DRAM**
9. **MultiThread**
   - **Threaded processors**
   - **SMT, DMT**
   - **Multiscalar**

COURSE OUTLINE (cont)

10. Future CPU challenges

- **Integer vs special processing**
- **Single core vs multiple**
- **The S curve in performance/ solutions?**
- **More speculation?**

11. Streaming Data and Processing

- **Why streaming data, characteristics,**
- **Solutions space for streaming processing**
- **DSP as a co-processor**

12. Conclusion lecture

- **Discussion**

Uniqueness of VLSI Architecture Design

- **Related to Multidisciplinary process**
  - **Optics**
  - **Process**
  - **System**
  - **Package**
  - **Testing**
- **Components - part of a system**
- **Complexity**
- **Competitive market**
- **Risk/unknown**
- **Zero error**
  - **Debug**
  - **Bug Fix**
- **Cost/performance/time to market**

"[In the beginning] we had little idea of what we had started. ...I remember... saying, 'Okay, we've done integrated circuit. What do we do next?"

Gordon E. Moore

TRENDS IN VLSI

Sources:
Shekhar Borkar
Uri Weiser

The Iron Laws of Computer Architecture

- **For a given instructions stream**
  Performance = Frequency x Insts-per-cycle (IPC)
- **For a given task**
  Performance = Frequency x IPC / #inst
- **Power**
  Power = Active Power + Leakage
  Power = αCV^2f + F(C, V, T)
  (α: Activity, C: capacity, V: voltage, f: frequency, T: temperature)

Performance History

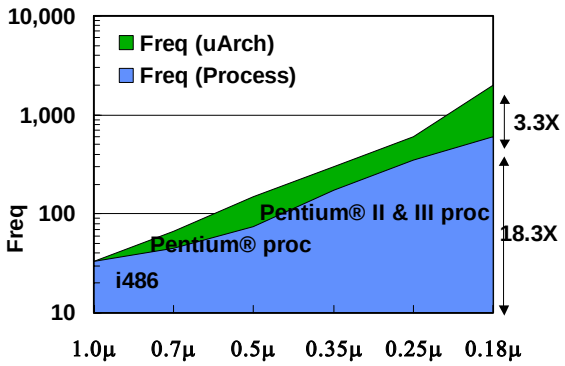

1.0u-0.18u, 1989-2001

Frequency increased 61X

1. **18.3X due to process technology**
2. **Additional 3.3X due to uArch**

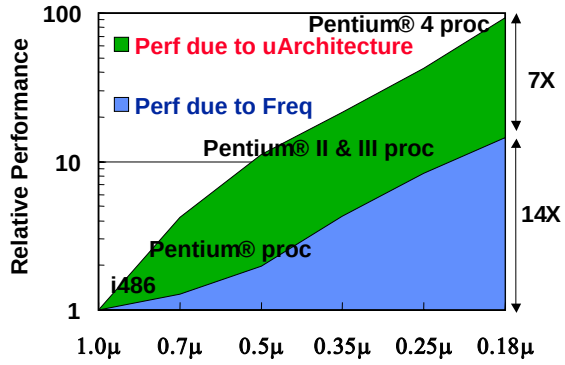

Performance increased ~100X

1. **14X due to process tech**
2. **Additional 7X due to uArch & design**

Process Technology: Minimum Feature Size

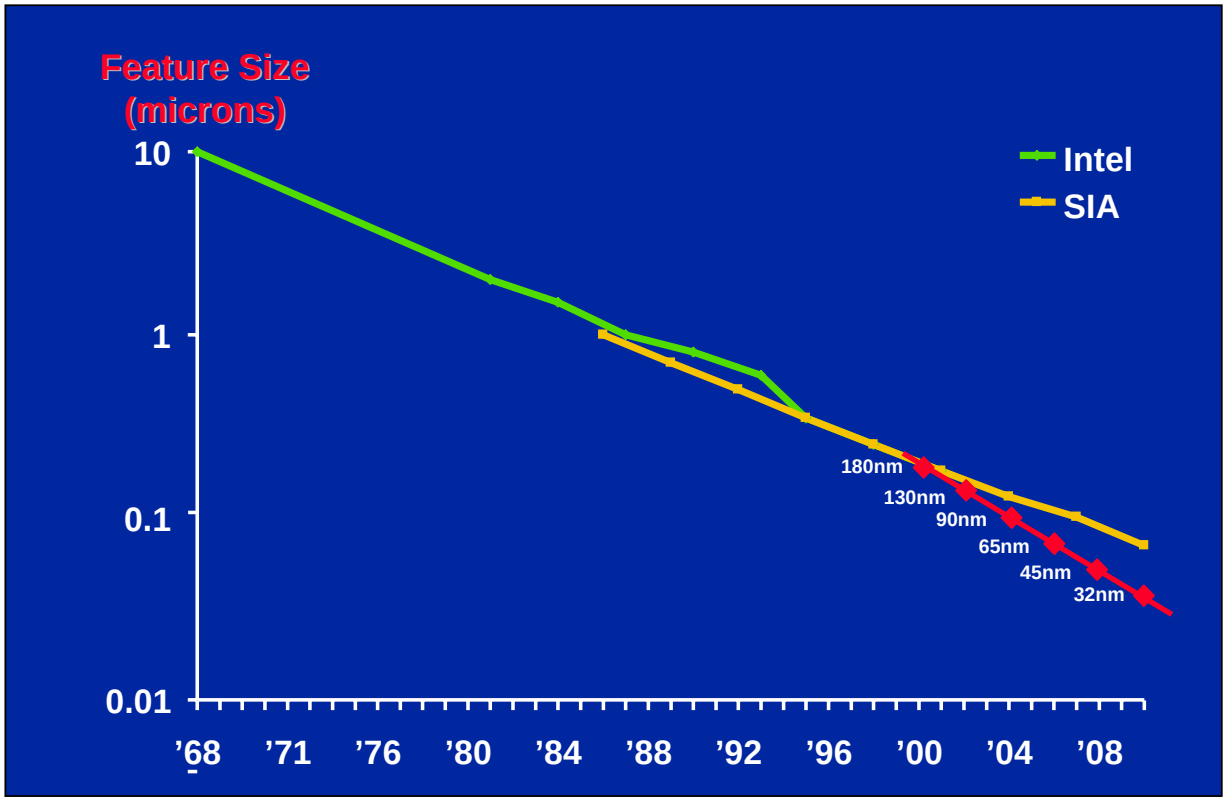

Source: Intel, SIA Technology Roadmap

Transistors on a Chip

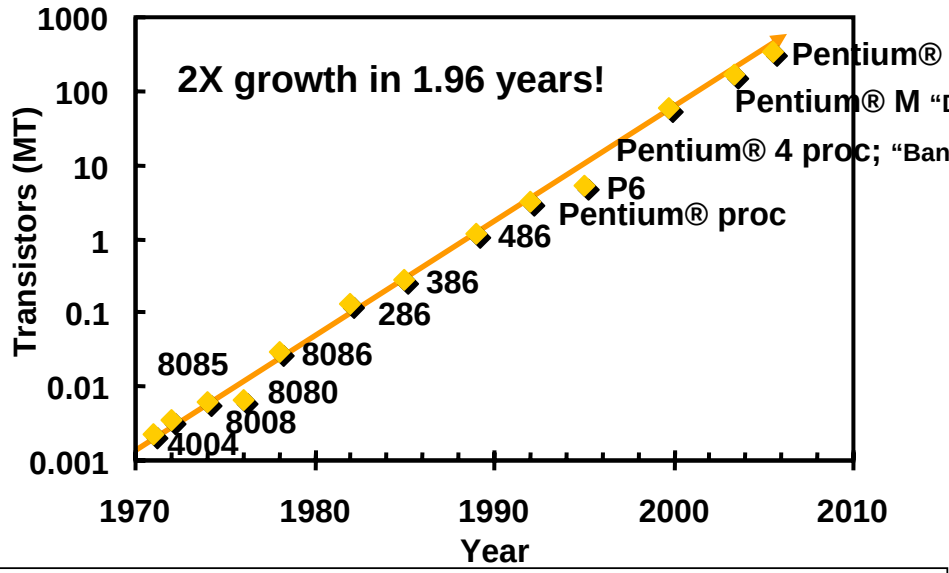

Transistors on a chip doubled every two years

Die Size Growth

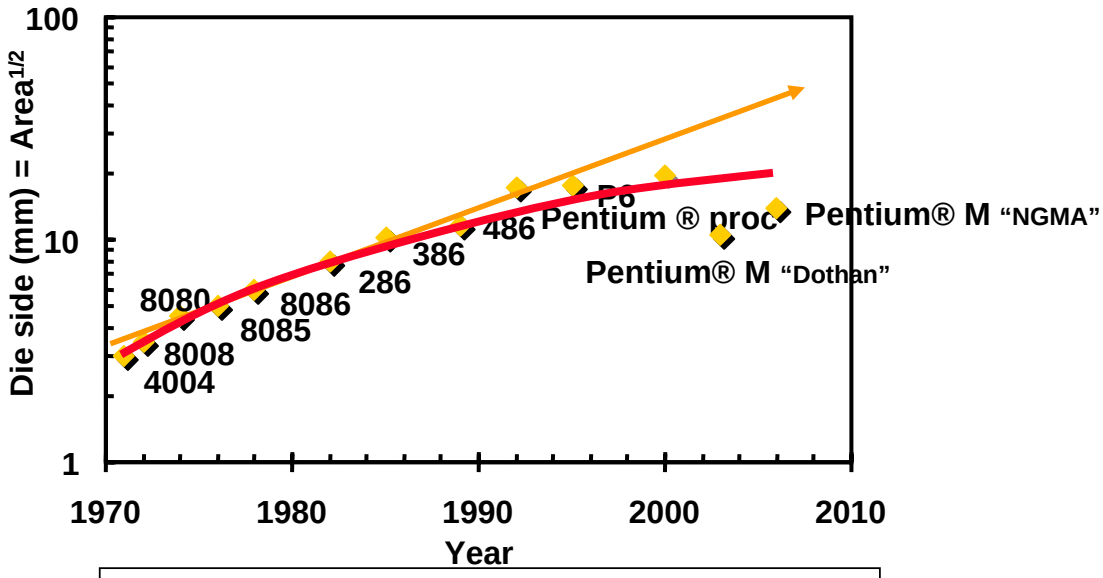

Die size grows? Is it saturated?

Frequency

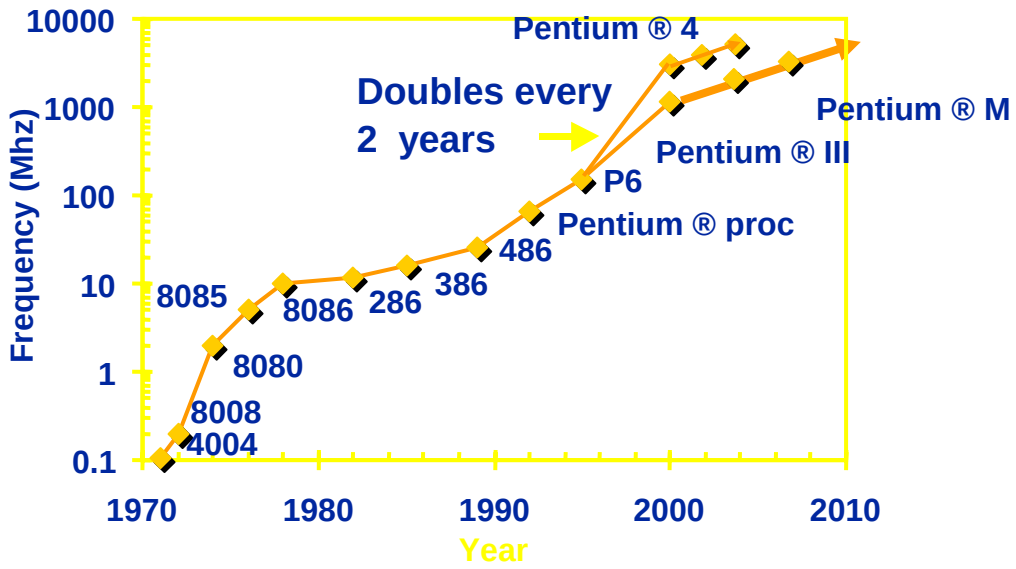

Lead Microprocessors frequency used to be doubled every 2 years
Not any more…

Frequency of Operation

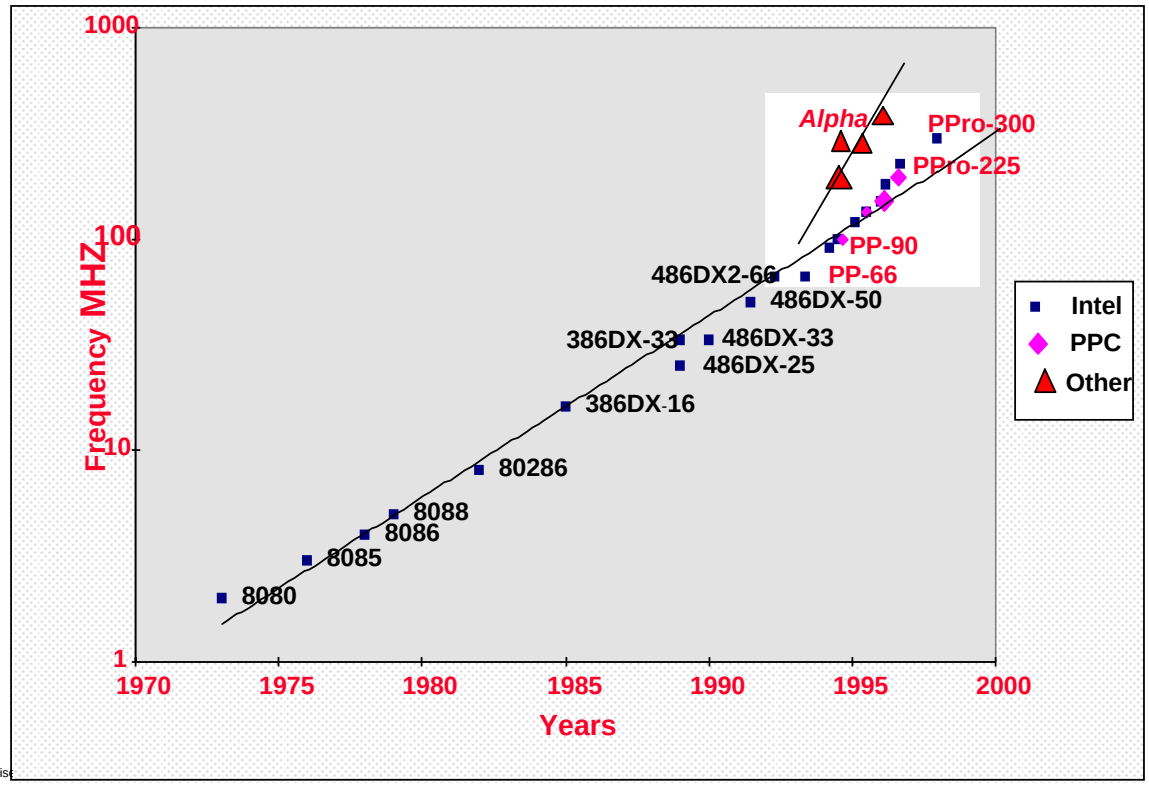

Uri Weiser

Frequency of Operation (cont.)

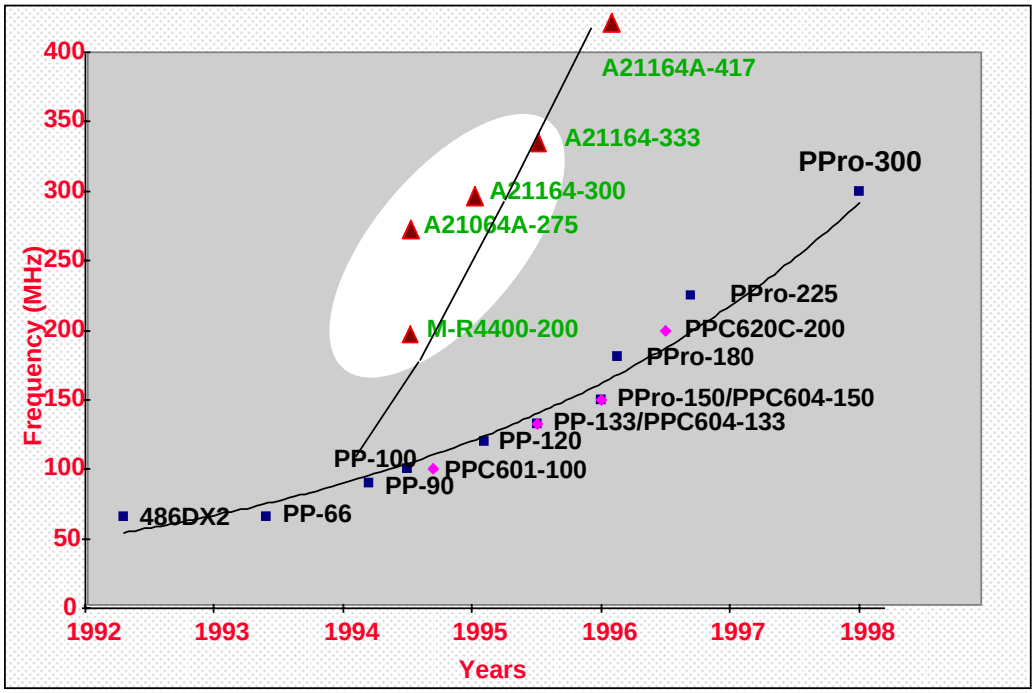

Brainiacs and Speed demons

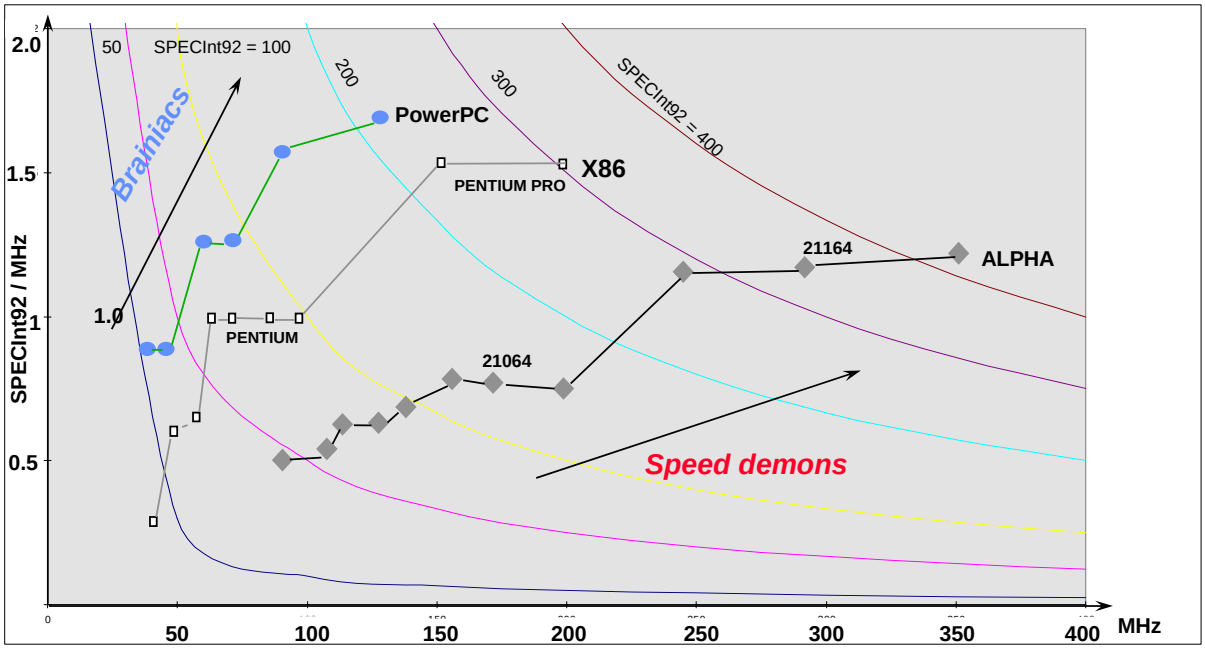

Uri Weiser – Ronny Ronen /VLSI_01_2006.ppt

Source: ISCA 95, p. 174

Trends of Future ProcessorsZ

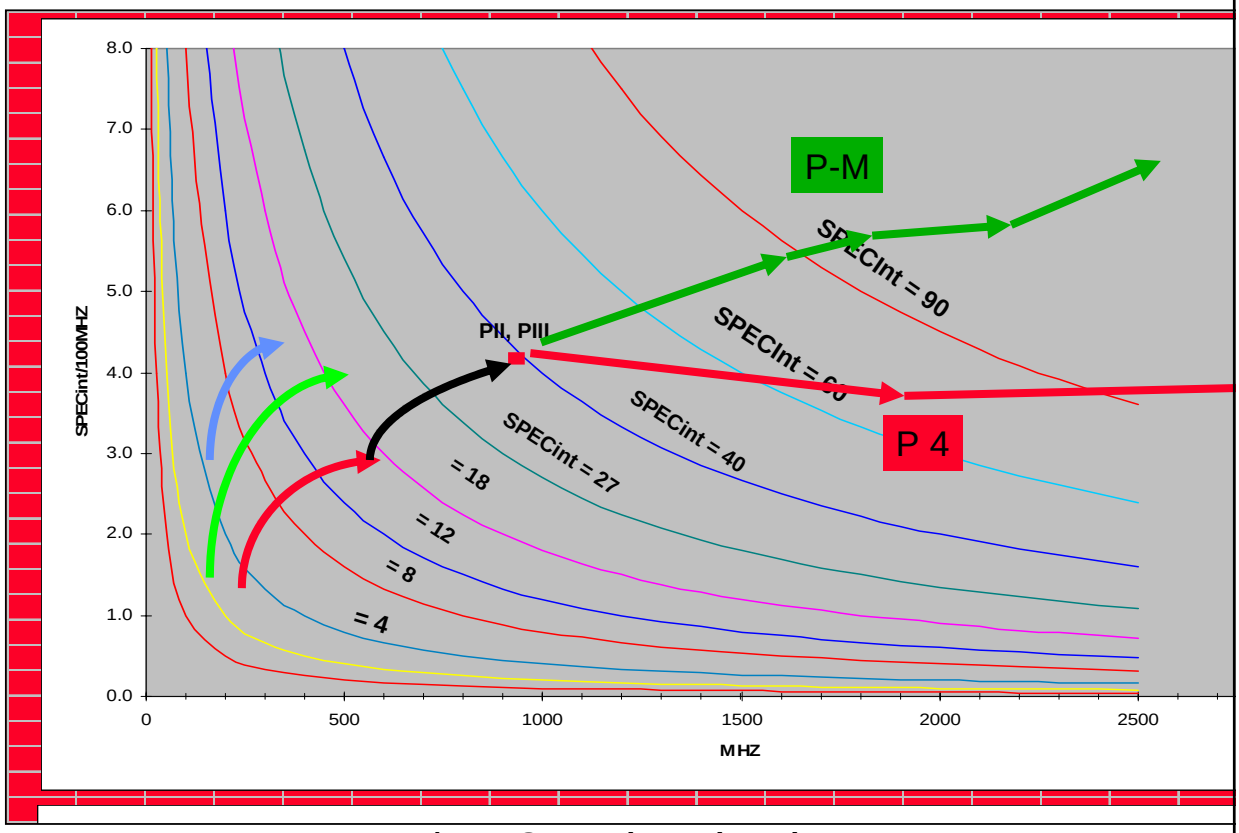

P4/PM - General trends only

Power

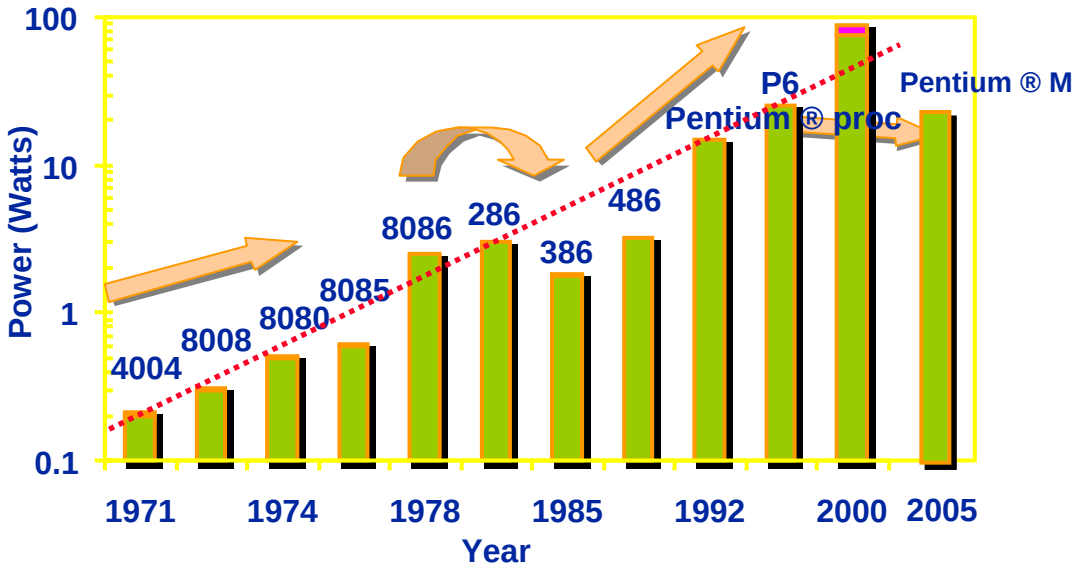

Lead Microprocessors power increased exponentially but the trend has changed…

Following Moore's Law

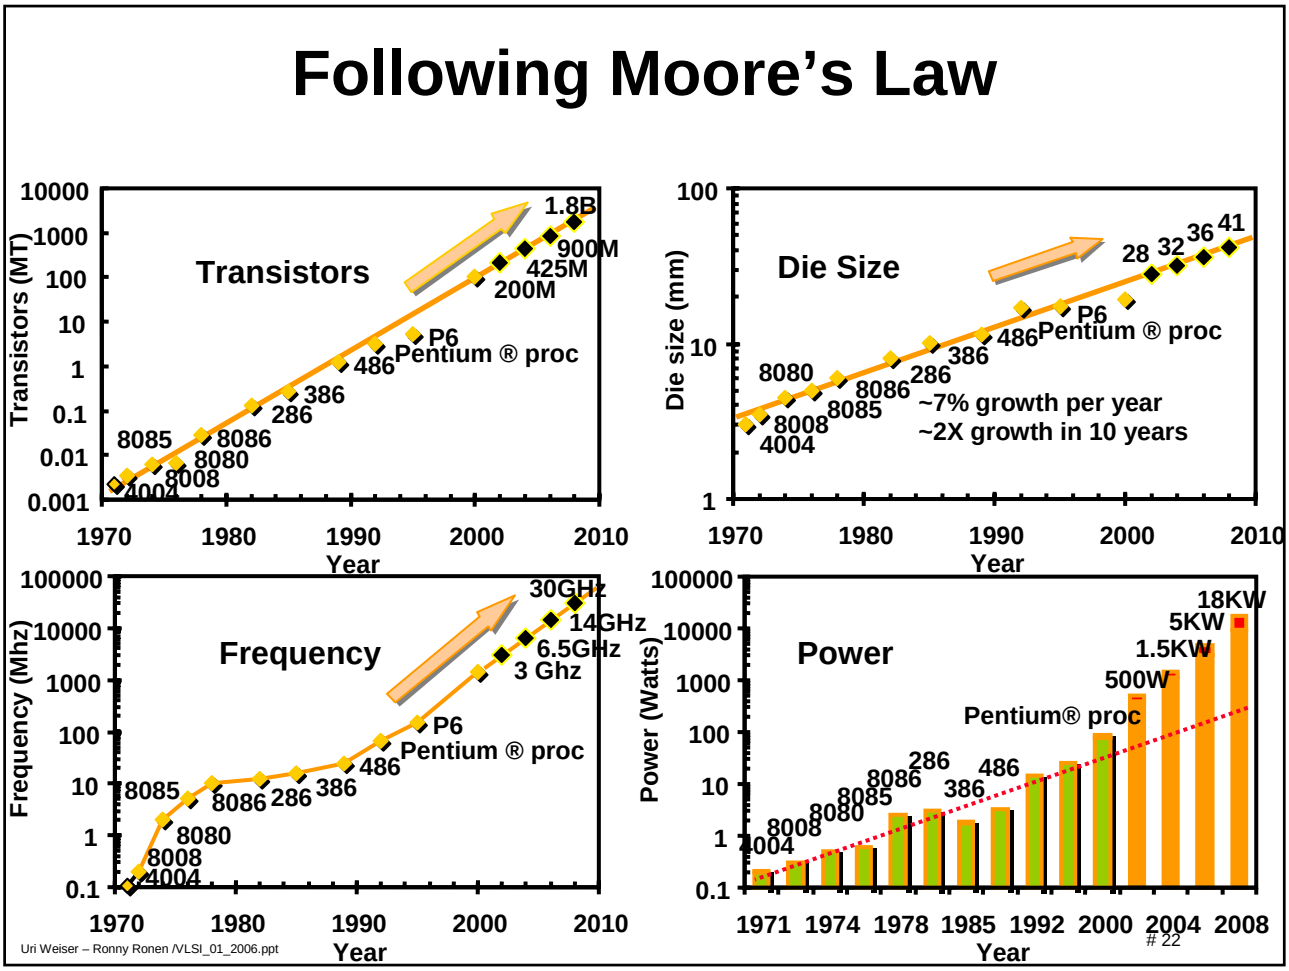

Logic Transistor Density

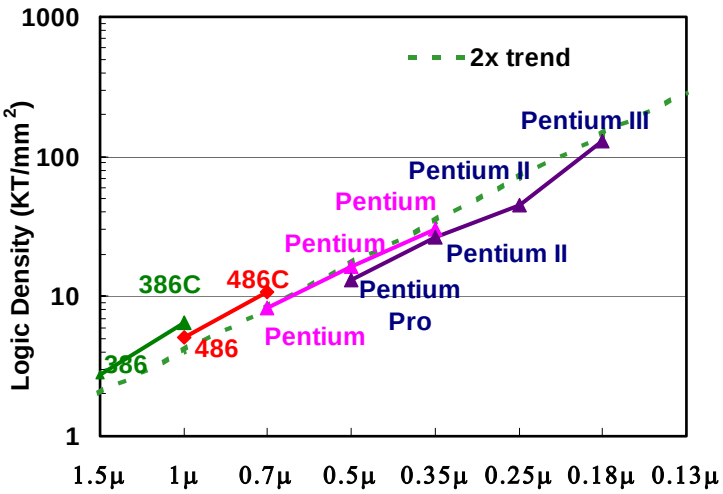

Shrinks & Compactions meet density goals
New u-Architectures drop density

Power Dissipation of Compactions

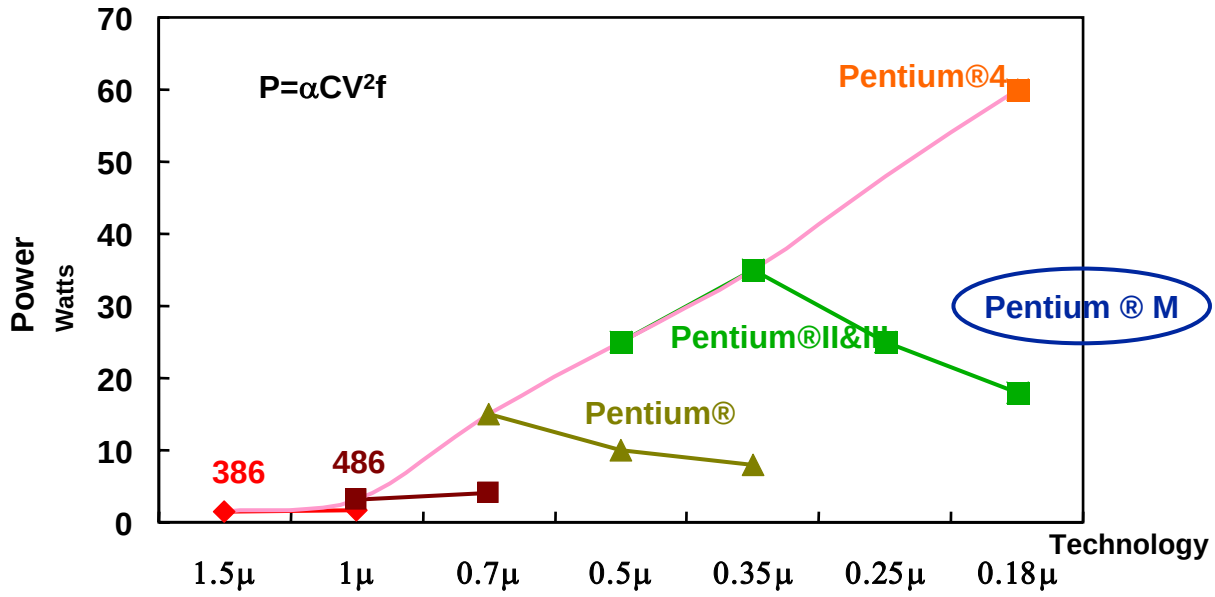

Lead processor power increases
Compactions provide higher performance at lower power

Active Power Trend

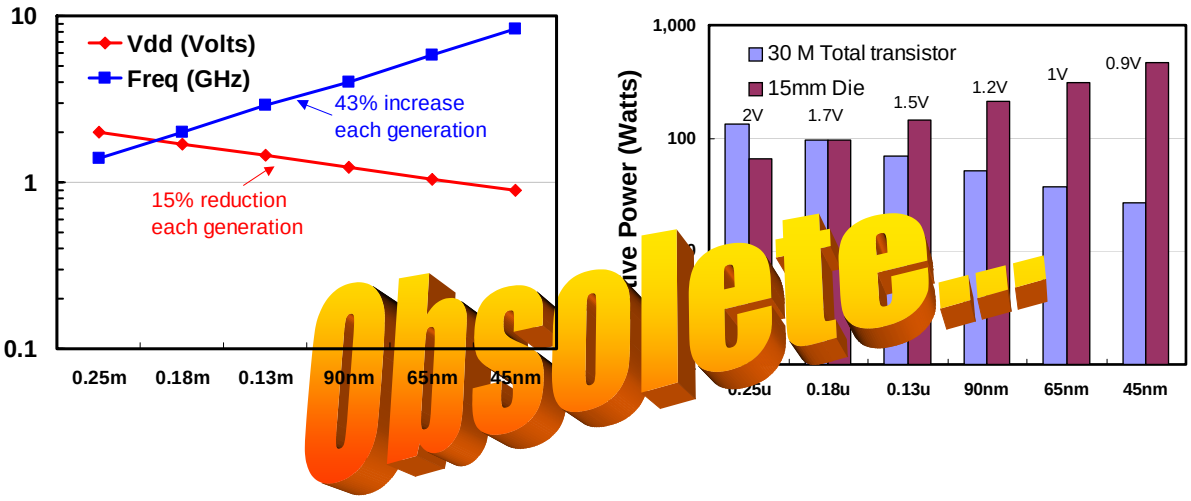

Obsolete...!

Must push thermal envelope beyond 300W
Else, die size must be reduced

Power density continues to get worse

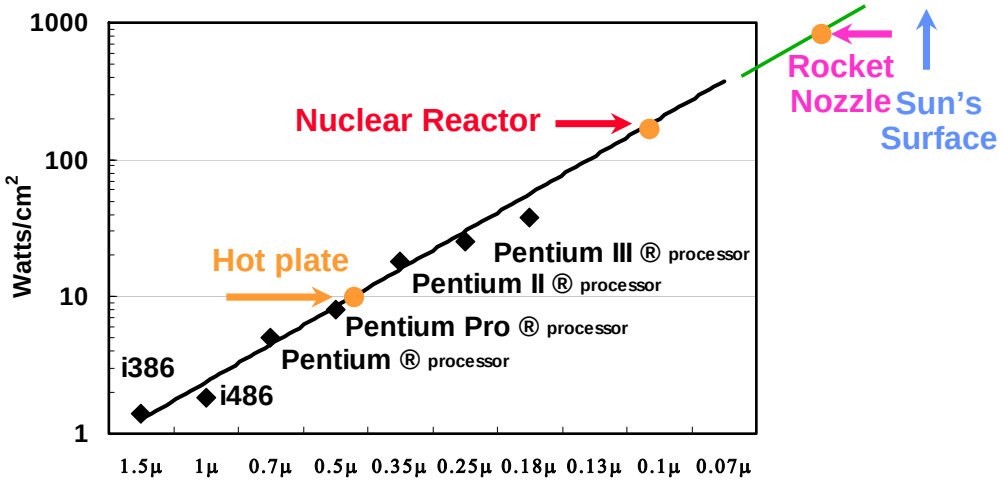

On Die Cache Memory

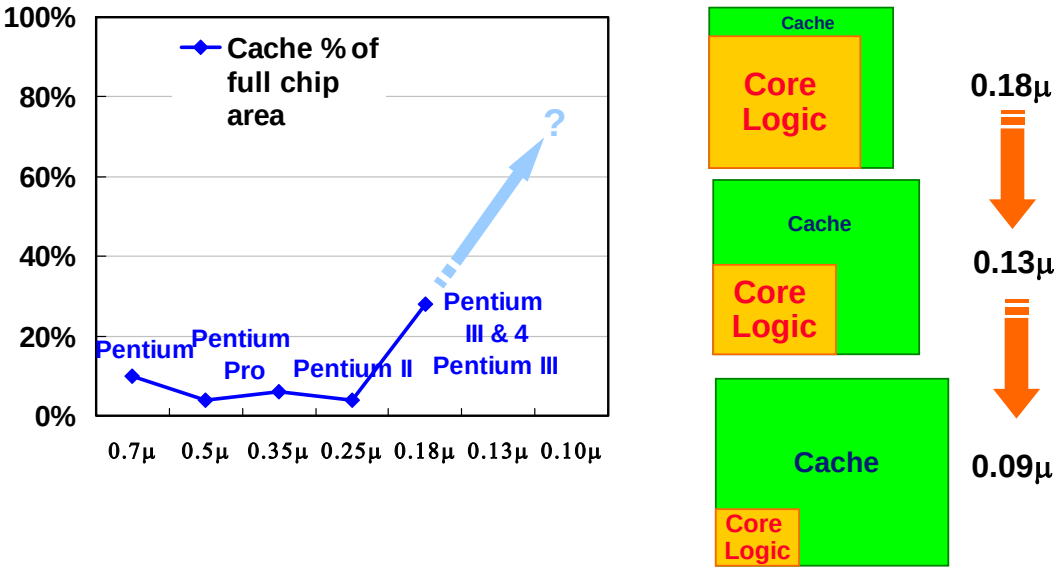

Larger % of die area will be memory

130 nm Generation Transistors

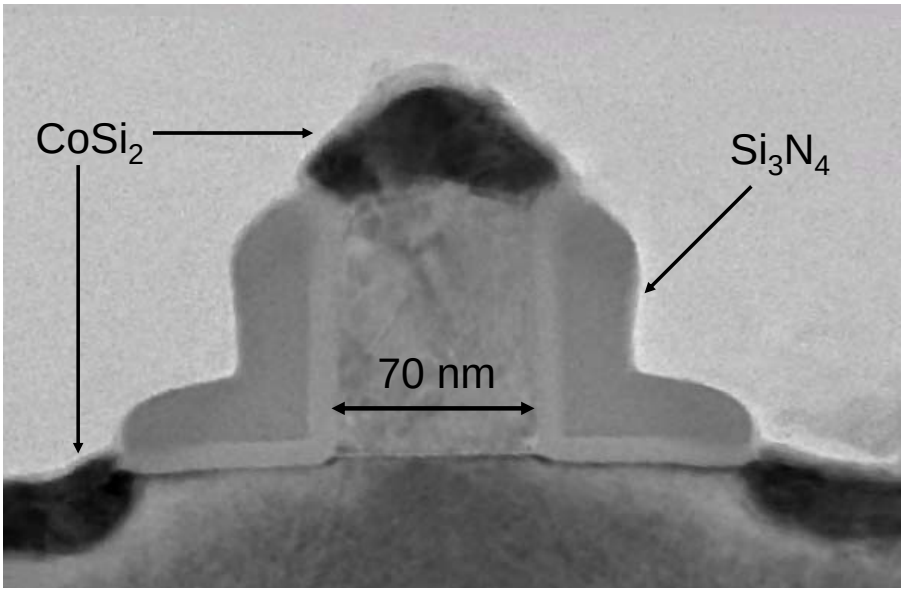

Source: Mark Bohr, Intel

Microprocessor Design Effort:

	1975	1980	1989	1992	1995	1998	2002
Designers	5	10	40	100	200	300	400
Layout	5	10	20	30	50	200	300
System Validation	--	--	5	20	30	50	70
Architects	1	1	1	10	15	20	40
~ Design Cost	$4M		$30M	→	$100M		$350

Trends:

- **Improvement in productivity**
- **Definition is more complex**

Two main trends

- **Process technology is not supporting power reduction/constant area any more**
- **Microarchitecture (in the original sense) is NOT a remedy to the performance thirsty market**

Process trend – the theory (cont)

Performance driven era vs. Power aware era

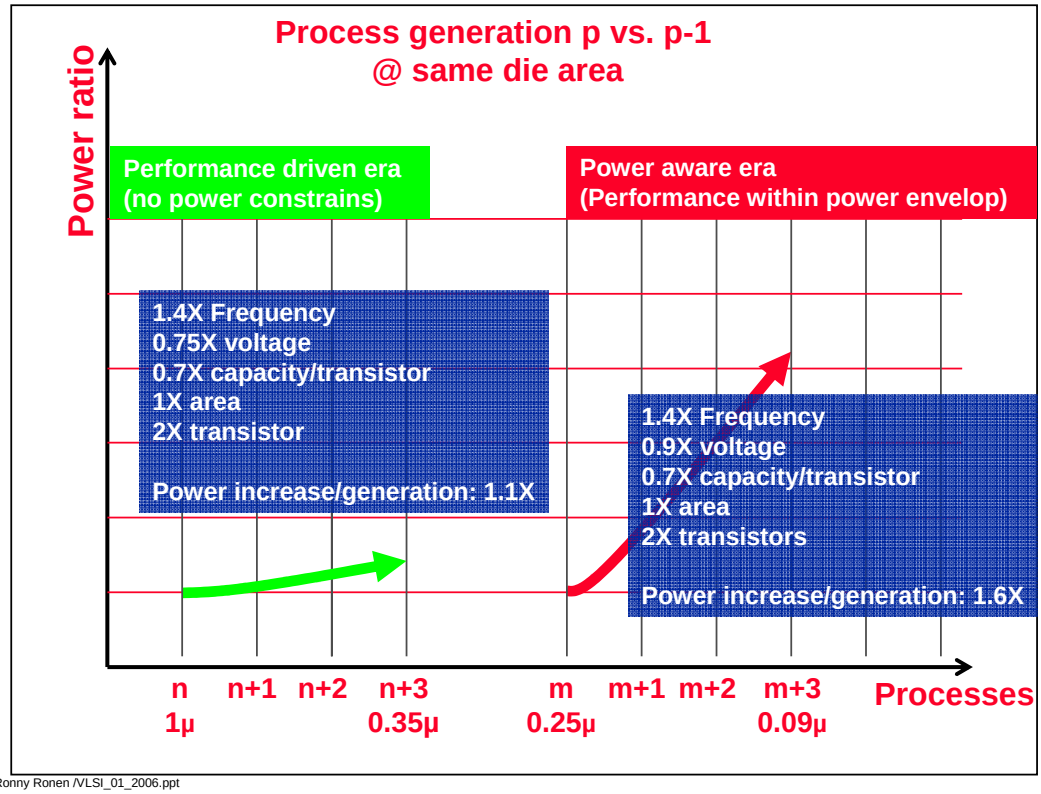

The power wall:

Processor roadmap trend – real life

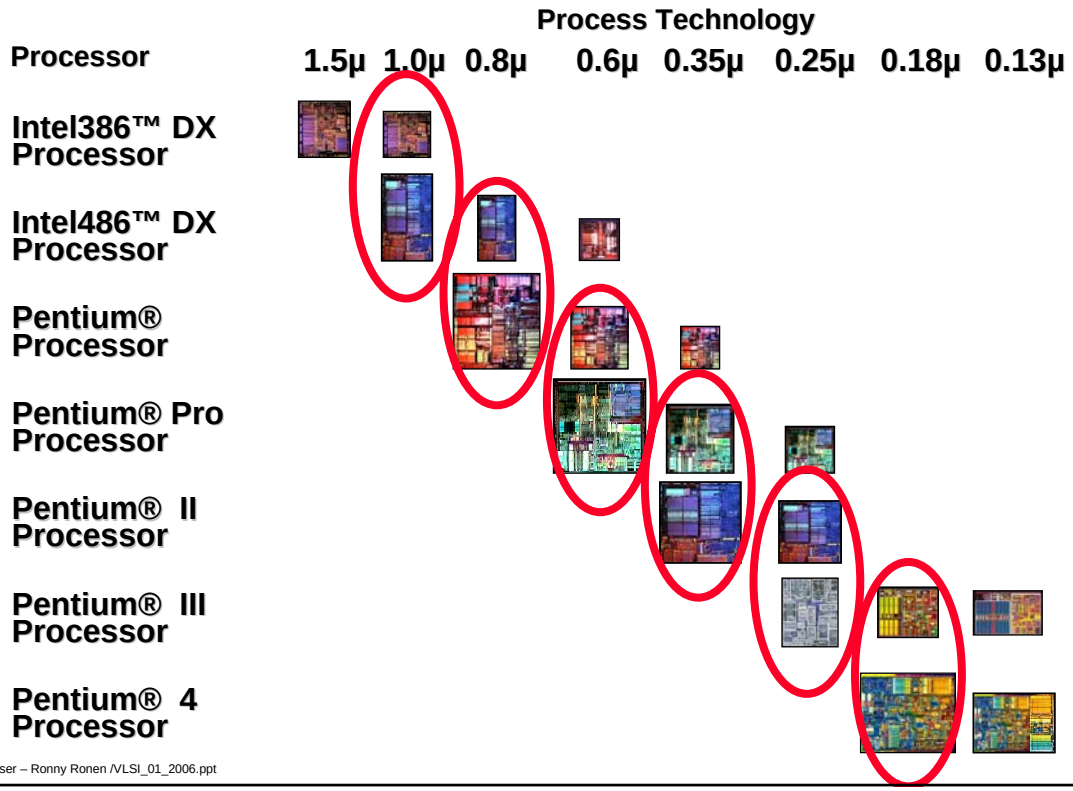

Uri Weiser – Ronny Ronen /VLSI_01_2006.ppt

32

Processor roadmap trend – real life (cont)

Extension of Pollack's Rule (Micro32, 1999)

Processor generation k vs. k-1 compacted @ the same process technology

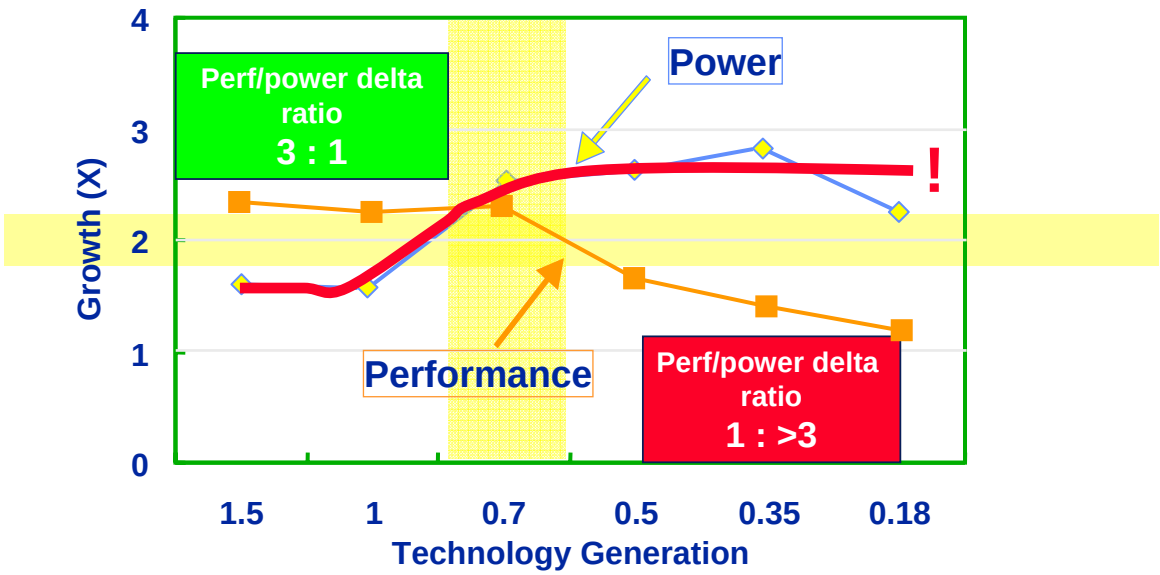

Multi-core

- Single core trends are gloomy
  - Performance is difficult to exploit
  - Perf/Power and Perf/Area deteriorate
- Fill the die w/ cache has limited value
  - Between 2MB and 4MB cache – gain is diminishing
- Solution: More cores on a die
  - Advantages – higher performance, higher perf/power, less complexity, lower wire delays
  - Challenges – scalability of interconnect, scalability of application, limited memory BW
- 1st incarnations
  - IBM Power
  - Intel's Pentium-D + Yonah
- Other solutions – more integration - memory controller, special functions, graphics, etc…

Better balance of cores & caches

Reliability

- Defects in processors
  - Soft Errors (Alpha particles, etc…)
  - Semi-Hard – voltage/frequency related
  - Hard – permanent failure
- Measured in FITs (failure in time)
  - 1 FIT - 1 failure in 10^9 hours
  - Requirements for systems – in the range of 1000s-10000s FITs
- Huge problem for servers (=many devices)
- Solutions for storage devices (Memory, caches)
  - Parity, ECC, etc…
- Solution for Logic - redundancy

Better balance of cores & caches

Summary

- More transistors
- Lower delay time

➔Increased Frequency

➔Increased Instruction Per Cycle

➔Increased Performance

➔Increased Power

➔Increased Error rate